

7N-63-TM
91482

P-6

N92-70679

Unclas
Z9/63 0091482

(NASA-TM-107887) FUNDAMENTAL RESEARCH IN
ARTIFICIAL INTELLIGENCE AT NASA (NASA) 6 p

Fundamental Research in Artificial Intelligence at NASA

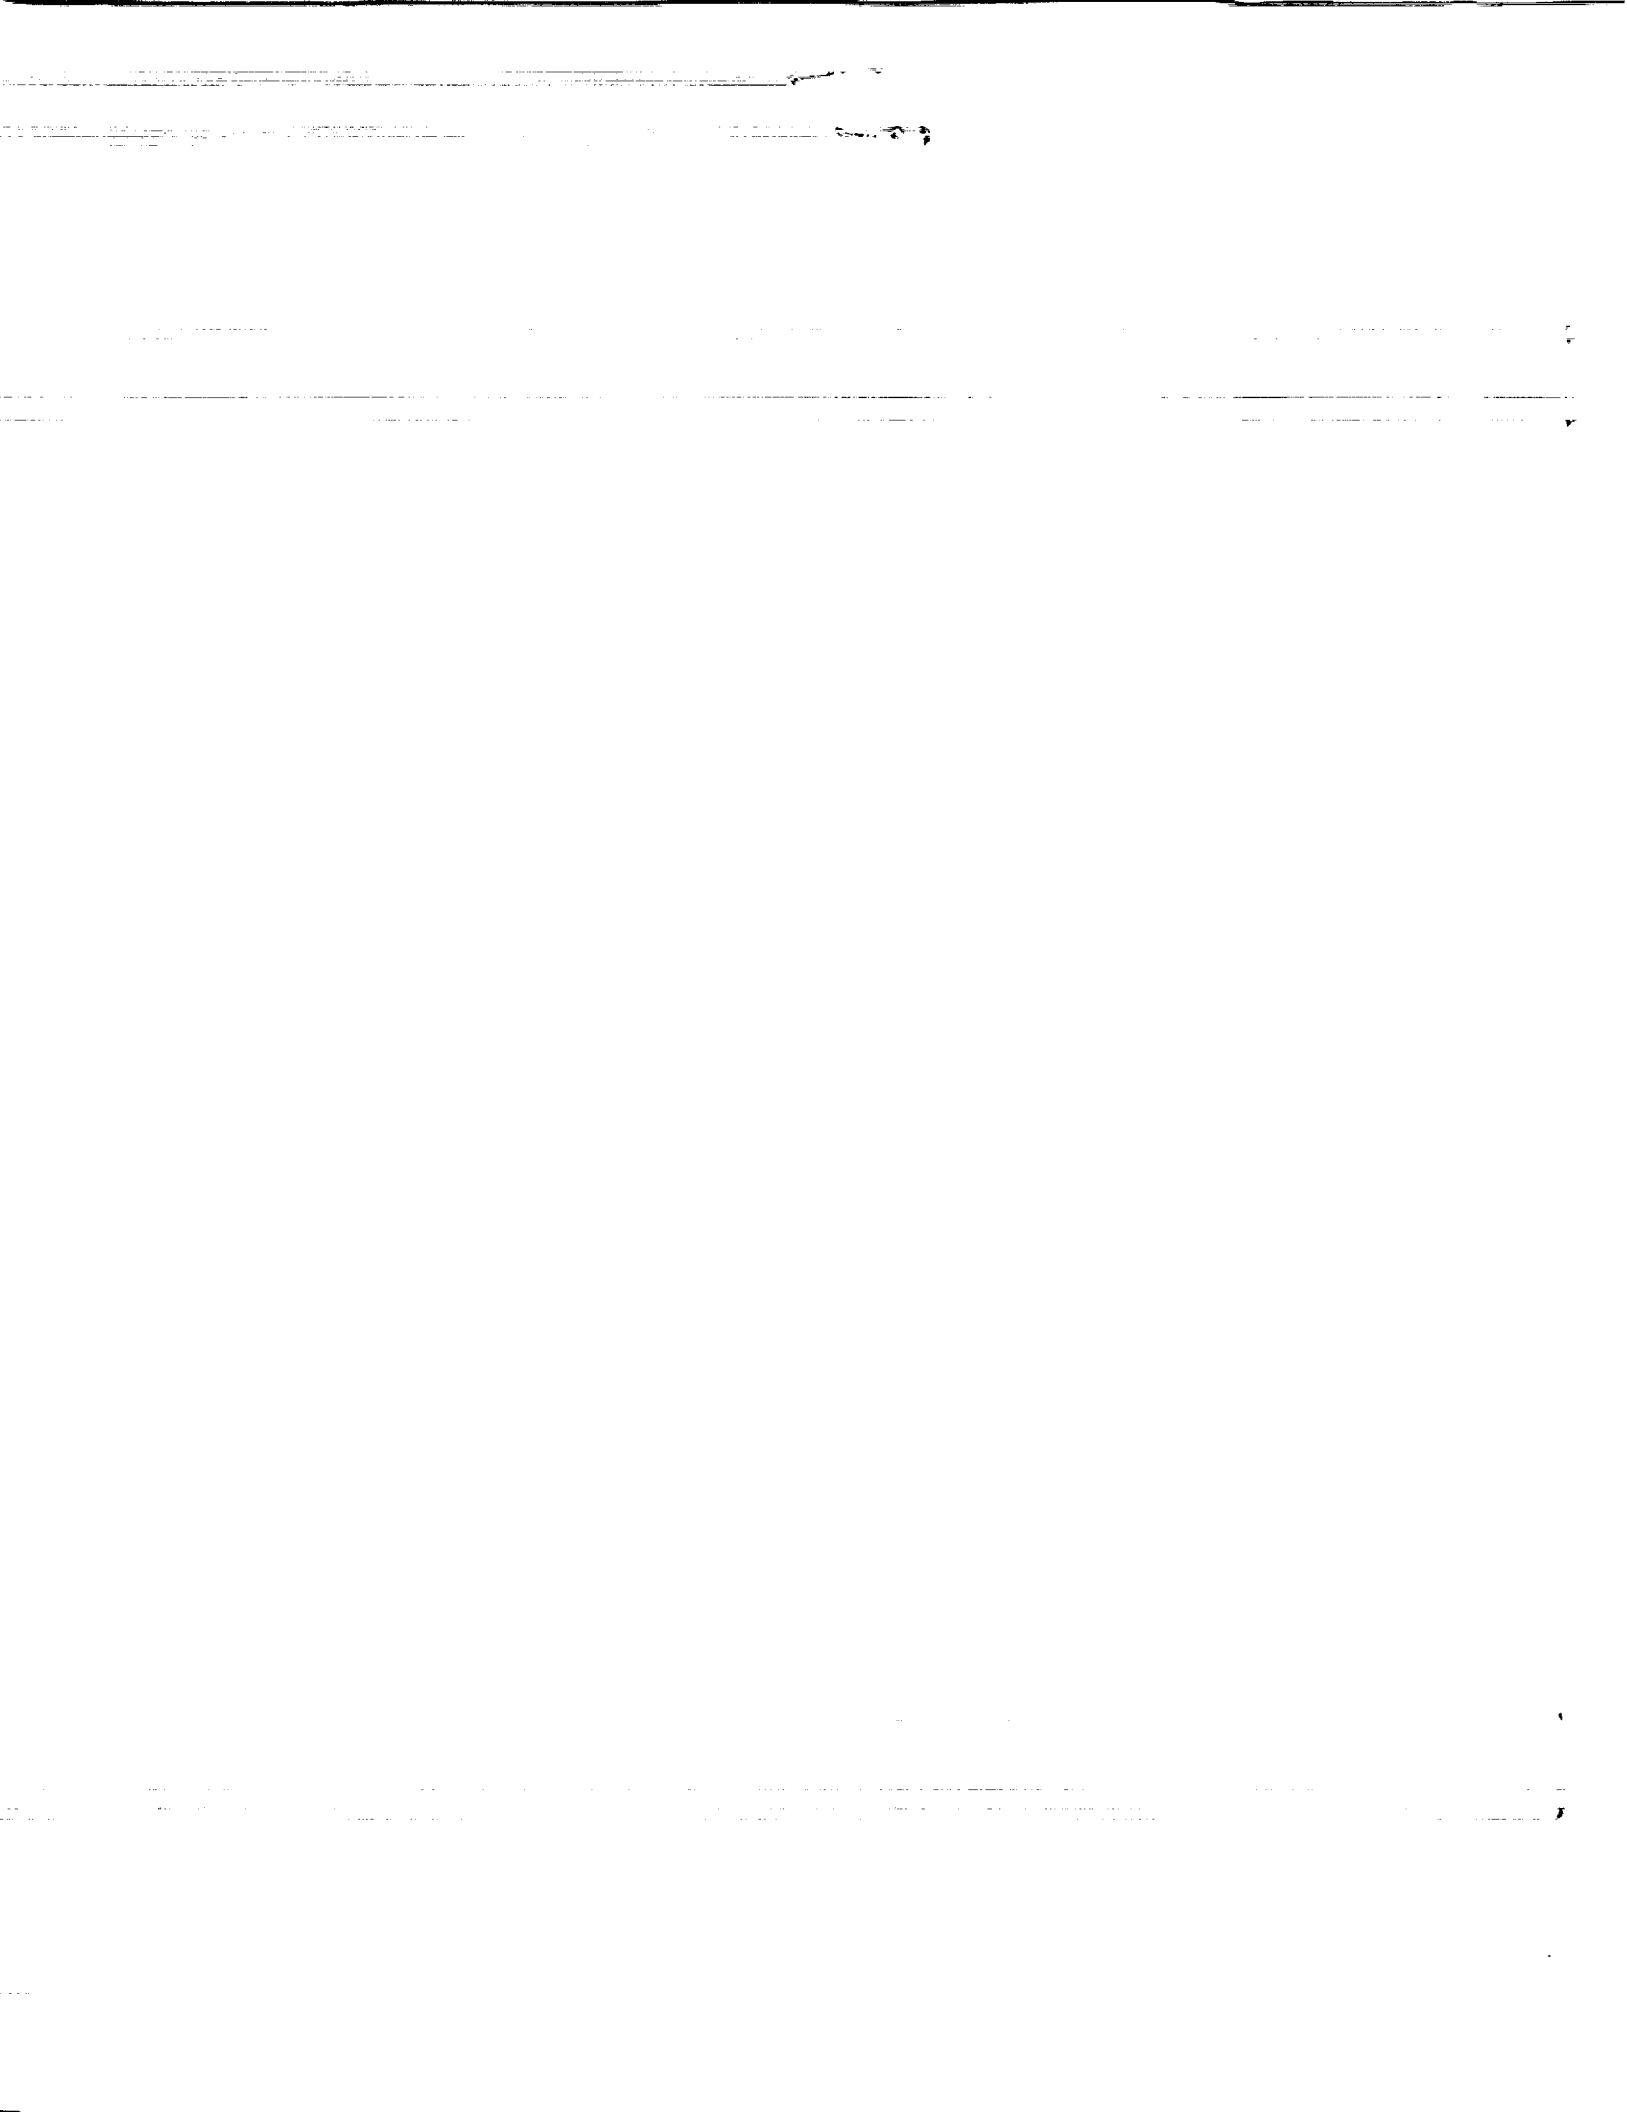
PETER FRIEDLAND

AI RESEARCH BRANCH, MAIL STOP 244-17
NASA AMES RESEARCH CENTER
MOFFETT FIELD, CA 94035

NASA Ames Research Center
Artificial Intelligence Research Branch

Technical Report FIA-90-11-01-02

November, 1990



Fundamental Research In Artificial Intelligence at NASA

**Peter Friedland
Artificial Intelligence Research Branch
Mail Stop 244-17
NASA Ames Research Center
Moffett Field, Ca. 94035**

415-604-4277 (office), 415-604-6997 (fax)

Abstract -- This paper describes basic research at NASA in the field of artificial intelligence. The work is conducted at the Ames Research Center and the Jet Propulsion Laboratory, primarily under the auspices of the NASA-wide Artificial Intelligence Program in the Office of Aeronautics, Exploration and Technology. The research is aimed at solving long-term NASA problems in missions operations, spacecraft autonomy, preservation of corporate knowledge about NASA missions and vehicles, and management/analysis of scientific and engineering data. From a scientific point of view, the research is broken into the categories of: planning and scheduling; machine learning; and design of and reasoning about large-scale physical systems.

Introduction

In 1987, NASA began an ambitious program in artificial intelligence research and development, aimed at producing solutions to the challenging variety of problems faced in current and future space missions. Those problems can be divided into four major classes:

- Intelligent Assistance for Mission Operations--the development of AI-based systems that can act as "intelligence amplifiers" for ground and space-based humans who have the responsibility to conduct mission operations

- Scientific and Engineering Data Analysis Tools--the development of tools to assist in the analysis of vast amounts of science and engineering data resulting from NASA missions

- Onboard Systems for Diagnosis, Planning, and Intelligent Control--the development of real-time, in-the-loop systems for planning, control, diagnosis, and fault correction of current and future spacecraft

- Capture, Integration, and Preservation of Life-Cycle Knowledge--the development of mechanisms for acquiring, combining, maintaining, and utilizing knowledge relating to the devices NASA designs, builds, and operates over long life-spans.

The Artificial Intelligence Program (managed by Dr. Mel Montemerlo of the Information Science and Human Factors Division of the Office of Aeronautics, Exploration and Technology) supports many projects whose purpose is to produce short-term applications of AI to current Agency problems that fall into these classes. This work tends to occur at the NASA Centers which support major current missions (for example, Johnson and Kennedy Space Centers for Shuttle launch and mission operations; Marshall Space Flight Center and Lewis Research Center for power systems applications; and the Jet Propulsion Laboratory and Goddard Space Flight Center for unmanned mission operations). In addition, the AI Program includes a major effort in basic AI research, internally through a sizable laboratory at NASA Ames Research Center and several projects at the Jet Propulsion Laboratory, and externally in

the form of grants to academic laboratories.

Although almost all subfields of AI have strong relevance to the mission problems described above, three broad topics have been chosen to provide focus to the basic research program. These are:

- Planning and Scheduling--deciding on a sequence of actions to achieve a set of complex goals and determining how to allocate resources to carry out those actions

- Machine Learning--techniques for forming theories about natural and man-made phenomena, and for improving the problem-solving performance of computational systems over time

- Design of and Reasoning about Large-Scale Physical Systems--research on knowledge acquisition, knowledge representation, combination of knowledge from many different sources, and multi-purpose utilization of knowledge bases large enough to represent NASA devices like the Hubble Space Telescope and Space Station Freedom.

The remainder of this paper will briefly describe much of the ongoing in-house research at NASA in these three areas of AI. Major projects and project leaders will be identified; unless otherwise specified, all work described is being conducted in the Artificial Intelligence Research Branch of NASA Ames Research Center.

Planning and Scheduling

- Constraint-based, iterative improvement methods for scheduling (Monte Zweben)--this project covers two important aspects of scheduling research. The first is the development of a constraint management system able to represent and manipulate a wide variety

of temporal relations among activities. The second involves the use of methods which are guaranteed to produce a usable schedule almost immediately and then incrementally improve it by use of simulated annealing (these are known as "anytime" methods). The research has already led to a functional scheduling tool now being tested on Shuttle Orbiter processing.

- Repair methods for scheduling (Eric Biefeld-JPL)--research on the use of heuristic techniques for improving schedules with oversubscribed resources by intelligent "shuffling" of tasks. A particular focus here is on interactive methods that involve human cooperation in schedule improvement.

- Multi-agent planning (Amy Lansky)--this work focuses on the problem of generating multi-agent plans for domains with complex coordination requirements. It deals with both action generation and action ordering, as well as scheduling issues such as resource allocation and timing. The concept of "locality," i.e. the determination of which goals and agents are affected by any given action in order to limit search during planning, is a critical part of the research.

- Integrating planning, scheduling, and control (Mark Drummond)--research in this area, sometimes known as "intelligent agents" has a central theme that planning is not complete upon generation of an initial plan. Rather, it is recognized that the actions in a plan, when executed in a real environment, can easily make changes in the environment that should cause modifications in the original plan. Therefore, the systems being built in this project are able to continuously monitor plan execution, note significant changes caused either by plan actions or by external forces, and decide to replan when necessary.

- Navigation planning for mobile robots (Dave Miller-JPL)--work on techniques for efficient path planning of

mobile robots (such as a planetary rover) and on integrating those methods with low-level robotic control algorithms. This research is also considering the interaction of navigation planning with scientific operations planning as well as the coordination of significant numbers of small robots.

Machine Learning

- Bayesian learning (Peter Cheeseman)--research on the application of Bayesian probability theory to the classification of large, potentially noisy, databases. The current technique is able to efficiently find classes in both real-valued and discrete data with no prior information on the number of classes present. This research has already resulted in a new catalog of infrared astronomical sources that was published as a NASA Reference Publication.

- Efficient learning algorithms (Phil Laird)--the development of methods for the production of software able to efficiently and reliably monitor complex spaceborne devices and react to changing conditions by learning. Current research is on unsupervised learning algorithms that can understand, by analysis of sensor values, when devices are beginning to behave "abnormally" and can both inform human operators and modify control programs automatically.

- Theory formation (Michael Sims)--research on initial formation and subsequent modification of theories about how the natural world "operates." This work includes overall modeling of the theory formation process as well as the utilization of theories to discover new concepts relevant to the system being studied.

- An integrated architecture for learning (Pat Langley)--the development of a comprehensive learning architecture for an agent that can interact with an unpredictable environment and acquire

knowledge from its experience. The architecture currently being studied uses a hierarchy of probabilistic concepts to store knowledge about physical objects, plans, and motor control schemas. It uses a process of heuristic classification to add the knowledge of new experiences to the hierarchy.

- Learning and performance improvement in scheduling (Steve Minton)--the integration of explanation-based learning methods with scheduling systems in order to produce schedulers that improve their performance over time. This involves two areas of research: the design of a framework for scheduling heuristics and the development of mechanisms for automatically learning such heuristics during operation of a scheduler.

- Learning in diagnosis (Deepak Kulkarni)--the use of concept formation methods to improve diagnosis of faults in complex devices. Current research focuses on the formulation of active experiments that are able to discriminate among competing fault hypotheses. The goal is to either automatically, or through human action, execute such an experiment, and, by observing device behavior, prune the active list of hypotheses.

Design of and Reasoning about Large-Scale Physical Systems

- Knowledge acquisition during design (Catherine Baudin)--the goal of this research is to address the problem of lost information through the lifecycle of an engineered device by developing an "electronic designer's notebook" that can "painlessly" record all information about alternative designs and both extract (from the human designer) and induce the rationale behind design decisions. In addition, there is work in progress on analyzing designs to determine how alternative designs meet device requirements.

methods like hypermedia and virtual reality devices.

- Knowledge compilation (Rich Keller)--this work explores the possibility of automatically producing efficient expert systems for specific tasks (e.g. diagnosis and re-design) from a large-scale, generic knowledge base containing structural and functional information about an engineered device. If successful, this approach should greatly facilitate both the updating and the validation process for expert systems since it is far easier and more verifiable to make changes to an underlying device model than to hunt down, test, and modify specific associational rules.

- Selective monitoring (Rich Doyle-JPL)--the goal of this project is to improve the process of monitoring complex devices using large collections of sensors. The research includes building a model of expected sensor values, interpreting actual data, and predicting the state of the physical system based on those values. Potentially, this will allow the dynamic determination of which sensors, out of a huge number for complicated devices like Space Station Freedom, are most relevant to any given problem-solving context.

- Approximate control (Hamid Berenji)--research on the integration of conventional control methods (like state feedback control) with symbolic methods (particularly those developed from fuzzy logic techniques). A major goal of this work is to be able to effectively utilize both quantitative and qualitative knowledge about a device to reason effectively in uncertain or rapidly changing situations.

- Computer-integrated documentation (Guy Boy)--research aimed at greatly improving human retrieval of information from the vast array of written documents relevant to any particular space mission. The work has two major foci: understanding the role of context and user preferences in retrieving information, and integrating AI-based knowledge manipulation techniques with advanced interaction

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

| | | | | | |
|---|---|--|----------------------------|---|--|
| 1. AGENCY USE ONLY (Leave blank) | | 2. REPORT DATE Dates attached | | 3. REPORT TYPE AND DATES COVERED | |
| 4. TITLE AND SUBTITLE Titles/Authors - Attached | | | | 5. FUNDING NUMBERS | |
| 6. AUTHOR(S) | | | | | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Code FIA - Artificial Intelligence Research Branch Information Sciences Division | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER Attached | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Nasa/Ames Research Center Moffett Field, CA. 94035-1000 | | | | 10. SPONSORING/MONITORING AGENCY REPORT NUMBER | |
| 11. SUPPLEMENTARY NOTES | | | | | |
| 12a. DISTRIBUTION/AVAILABILITY STATEMENT Available for Public Distribution <i>Pete Fiedler</i> 5/14/92 BRANCH CHIEF | | | | 12b. DISTRIBUTION CODE | |
| 13. ABSTRACT (Maximum 200 words) Abstracts ATTACHED | | | | | |
| 14. SUBJECT TERMS | | | | 15. NUMBER OF PAGES | |
| | | | | 16. PRICE CODE | |
| 17. SECURITY CLASSIFICATION OF REPORT | 18. SECURITY CLASSIFICATION OF THIS PAGE | 19. SECURITY CLASSIFICATION OF ABSTRACT | 20. LIMITATION OF ABSTRACT | | |

FIA-90-09-24-01

Experimentation in Machine Discovery

DEEPAK KULKARNI AND HERBERT SIMON

September 1990

This report describes In this chapter we described KEKADA, a system that is capable of carrying out a complex series of experiments on problems from the history of science. The system incorporates a set of experimentation strategies that were extracted from the traces of the scientists' behaviors. It focuses on surprises to constrain its search, and uses its strategies to generate hypotheses and to carry out experiments. Some strategies are domain independent, whereas others incorporate knowledge of a specific domain. The domain-independent strategies include magnification, determining scope, divide and conquer, factor analysis, and relating different anomalous phenomena. KEKADA represents an experiment as a set of independent and dependent entities, with apparatus variables and a goal. It represents a theory either as a sequence of processes or as abstract hypotheses. This report describes KEKADA's response to a particular problem in biochemistry. On this and other problems, the system is capable of carrying out a complex series of experiments to refine domain theories. Analysis of the system and its behavior on a number of different problems has established its generality, but it has also revealed the reasons why the system, in its present form, would not be a good experimental scientist.

FIA-90-11-01-01

Planning and Scheduling Research at NASA Ames Research Center

PETER FRIEDLAND

November 1990

Planning and scheduling is the area of artificial intelligence research that focuses on the determination of a series of operations to achieve some set of (possibly) interacting goals and the placement of those operations in a timeline that allows them to be accomplished given available resources. This paper describes work in this area at the NASA Ames Research Center ranging from basic research in constraint-based reasoning and machine learning, to the development of efficient scheduling tools, to the application of such tools to complex Agency problems.

FIA-90-11-01-02

Fundamental Research in Artificial Intelligence at NASA

PETER FRIEDLAND

November 1990

This paper describes basic research at NASA in the field of artificial intelligence. The work is conducted at the Ames Research Center and the Jet Propulsion Laboratory, primarily under the auspices of the NASA-wide Artificial Intelligence Program in the office of Aeronautics, Exploration and Technology. The research is aimed at solving long-term NASA problems in missions operations, spacecraft autonomy, preservation of corporate knowledge about NASA missions and vehicles, and management/analysis of scientific and engineering data. From a scientific point of view, the research is broken into the categories of: planning and scheduling; machine learning; and design of and reasoning about large-scale physical systems.

